Autonomous Operations for the Crew Exploration Vehicle – Trade Study Design Considerations

James M. Crawford[1]and Charles R. Weisbin[2]

[1]NASA Ames Research Center, Moffett Field, CA 94035
[2]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109
[1](650) 604-1139, jc@email.arc.nasa.gov

Abstract. In the design of the operational concept for the CEV there are numerous choices regarding the locus of command and control. Systems such as power, propulsion, GN&C, life support, C&DH, etc. can be monitored and controlled by the flight crew, by onboard autonomous systems, by ground crew, or by ground autonomous systems. The decision of how to distribute control must be based on a complex trade-off between development and validation costs, operations costs, and reliability/risks. Getting these trade-offs wrong can lead to unnecessary growth in mission cost and risk, and unnecessary decreases in the time the crew has available for core exploration tasks. Over the next two years we will be performing an in depth analysis of the return on investment that we can expect from software tools that automate, or partially automate, the operation of CEV systems. In this paper we preview the issues in performing a trade study of this type, and the technical approach that will be used to gather and analyze the data required to perform this, and other similar, trade studies.

INTRODUCTION

The United States' space vision calls for a return to the moon by 2020 followed by manned missions to Mars and beyond. Unlike in the Apollo days, however, NASA's budget during this period will be nearly flat. This means that to ensure sustainability, significant attention will need to be paid to the lifecycle costs of space missions. Today's NASA budget includes approximately $7 billion for operation of the space station and space shuttle. The majority of the funds intended to cover the cost of the development of the Crew Exploration Vehicle (CEV) are scheduled to come from the phase out of the space shuttle. Once the first spirals of CEV development are complete, current plans call for continued development in parallel with lunar operations. However, with little or no new funds coming into NASA, if the operations costs of the CEV are similar to those of the space shuttle, then further CEV development spirals, which are essential for missions to Mars, will come to an end due to cost pressures once CEV operations on the moon begin.

There is no single "silver bullet" for reducing CEV lifecycle costs. Rather, there are a variety of operations costs that will need to be driven down. One such cost is routine operations. This includes such tasks as monitoring telemetry streams, planning crew activities, planning system and sub-system activities, performing root cause analysis for "routine" anomalies, etc. These tasks are interesting targets for partial or full automation because similar tasks have been automated in NASA's unmanned missions and/or in Defense Department projects. Further, for Mars exploration, most or all of these tasks will need to be done by the flight crew (due to communications latencies resulting from the finite speed of light). Thus automating routine operations has the potential to both reduce operations costs and enable crew self-sufficiency on Mars.

However, automation of decision making for the CEV is neither cheap nor risk free. Further, the CEV consists of many systems and sub-systems. For each of these systems and sub-systems one could, in principle, build automated systems for planning, sequencing, execution, monitoring, diagnosis, and prognosis. Thus there are tens to hundreds of possible autonomy applications. For each of these there would be significant costs in terms of software development, validation, training, and maintenance. It is unlikely that for each automated system the development and validation costs would outweigh the savings in operations costs. But it is very likely that for at least some of these automated systems the operations costs reductions would be much greater than the cost of developing and validating the autonomous system.

In principle determining which autonomy applications to build is straightforward. One simply compares the costs of system development and validation against the savings in operations costs. In practice there are several critical issues and complications:

- Operations costs are hard to quantify. The total costs of Station or Shuttle operations is known but the break down by system or sub-system is a challenge. Further, impacts of automation on operations costs are hard to predict or evaluate.
- Safety is hard to evaluate. Arguments can be made that automation will increase or decrease system safety. The need to ensure that safety is not compromised can prevent expected cost savings and/or increase development and validation costs for autonomous systems.
- Development and validation costs for autonomous systems are hard to estimate reliably.

Nevertheless, CEV operations costs must be reduced and decisions about applications of autonomy must be made. The purpose of this paper is to discuss design considerations in our ongoing effort to create a trade study to identify the autonomy applications which are most likely to significantly reduce lifecycle CEV costs. We discuss approaches to the challenges listed above. We conclude that while existing data can be used to indicate where CEV investment should be made, additional work is required to formalize processes to capture and analyze the data needed to make ongoing decisions about the role of autonomy in the exploration initiative.

AUTONOMY TRADE OFFS

Consider an autopilot for an airplane. If there is no autopilot then the pilot must continuously fly the airplane. If there is an autopilot then during the "boring" part of the trip the pilot can engage the autopilot and focus on other activities. This is a very simple case, but it illustrates all of the major issues in the automation of more complex systems:

- Validation: Safety could be compromised if the autopilot software is fault. There is, at least in principle, a danger that the autopilot could crash the airplane. The need for exhaustive testing and/or formal verification methods is a common theme in the use of autonomy for space mission as well.
- Response time: Safety is enhanced because the pilot is not being called on to perform a monotonous task. Similarly for space missions, autonomy can often be used to take error-prone tasks away from human operations. An additional safety issue in space flight is the need to respond to some anomalies more quickly than human reaction times (either because the required reaction times are extremely short or, for more complex cases that the flight crew has not been trained to handle, because of communication latencies to earth).
- Unexpected situations: Safety could be compromised if something unexpected happens while the autopilot is engaged – for example if another airplane appears which is on a collision course. Similarly, in space applications there may be conditions that were not anticipated by the designers and testers of the autonomous system that could endanger the crew and the mission.
- Development costs: Development costs are increased because of the need to develop and validate the autopilot. Similarly, development and validation costs for autonomous systems for space applications are substantial.
- Operations costs: Operations costs may be reduced. Obviously the airplane still needs a pilot, but without an autopilot long flights would require a team of pilots, and the number of flights a pilot could fly per week would be significantly reduced. Similarly autonomy can reduce the size of the ground team for space missions. Further, reductions in the size of the ground team, and simplifications in the tasks the ground team must do, have broad implications for training budgets, facilities, procedure development, etc.

These considerations can be summarized operationally in a decision diagram like the one shown in figure 1.

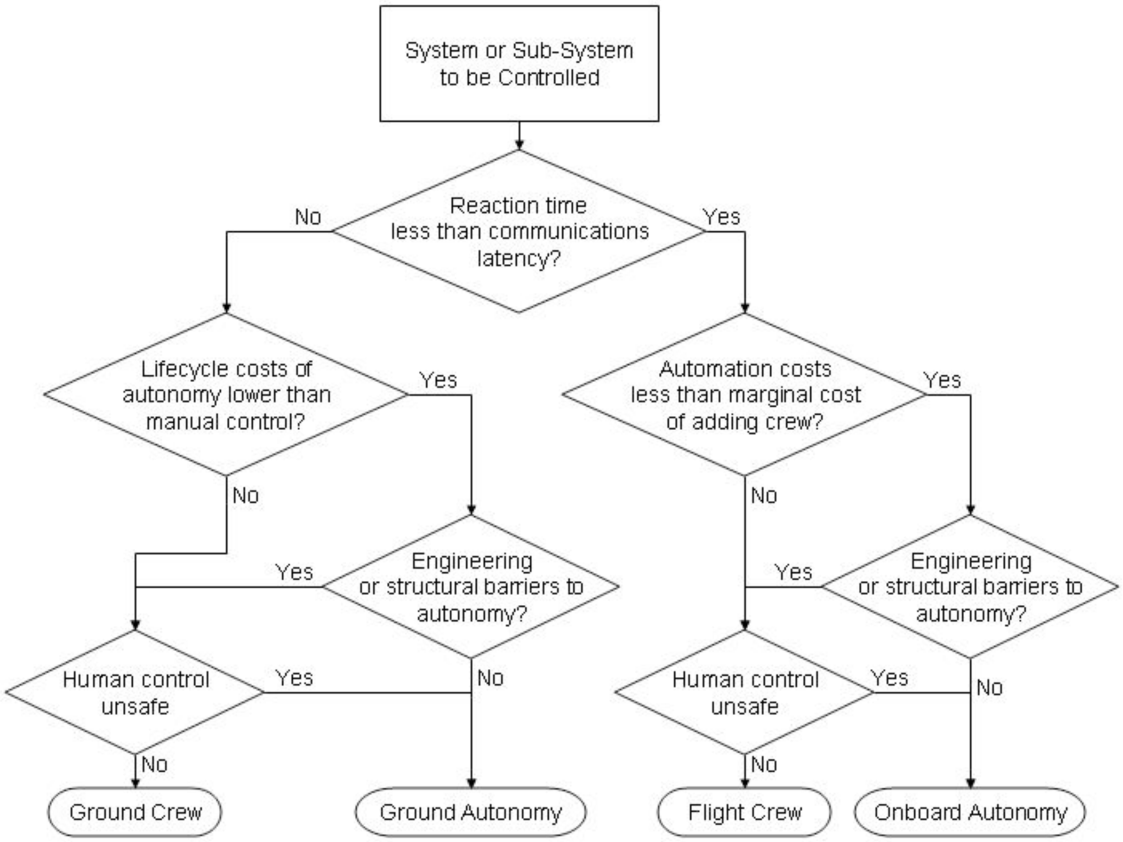

FIGURE 1. Autonomy Decision Diagram.

This decision diagram is a refinement of the general CEV goals of affordability, reliability/safety, and effectiveness. First priority is given to safety – putting control onboard or on the ground as driven by latency requirement and by the relative safety of human vs. autonomous control. When control can safely be manual or autonomous, the decision is driven by effectiveness and affordability – looking critically at lifecycle costs (development + validation + maintenance + training + operations). Figure 1 is a working hypothesis rather than a final conclusion but should illustrate the concept and motivate the data collection and analysis efforts discussed below.

The primary elements of this decision diagram are the following:

System or Sub-System to be Controlled: This could be either a major system such as power or propulsion, or a sub-system such as a part of a life support system. This analysis may have to be split temporally – e.g., ascent may require different reaction times than cruise, or increased autonomy may be required during crew sleep periods. Also, some systems have multiple classes of control decisions at different temporal granularities (e.g., life support). Interestingly, the same analysis also applies to systems of systems level decisions that can be made onboard or on the ground (e.g. mission level command and control).

Reaction time less than communications latency: Some control decisions, such as rocket firings during final docking maneuvers, must be made within fractions of a second. Others, such as control of biological life support, are made on time scales of weeks. Clearly if a decision must be made on a time scale less than the communication latencies from earth to the CEV then the decision must be made on board. In this case the tradeoff is between flight crew operations and automated systems. Conversely, if the decision can be made on earth then the tradeoff is between ground crew operations and automated systems. Note also that many more decisions must be made onboard for Martian operations than for lunar operations.

Lifecycle costs of autonomy lower than manual control: Stating this criterion is easy but collecting the requisite data is a challenge. Primary data required includes: the cost of the development of autonomous systems, the cost of validation and testing of autonomous systems, the relative cost of the development of efficient user interfaces for manual and autonomous systems, and the operations costs of the manual system.

Lifecycle costs of autonomous system less than marginal cost of adding flight crew: We assume for purposes of autonomy trade studies that other demands on flight crew time are fixed. Thus any crew time spent monitoring and controlling basic systems must require the addition of crewmembers (one could equivalently start from the assumption that the crew does all low latency monitoring and control, and measure the cost of flight autonomy against the savings resulting from decreasing the crew size). Adding additional flight crew members has major systems-of-systems impacts across the design on the CEV. Thus it is likely that this marginal cost would be quite high.

Engineering or Structural Barrier to Autonomy: In some cases, processes may be too complex for existing technology to automate, the mass-to-space constraints for the hardware may preclude the use of autonomy, or verification and validations techniques may not sufficient to bring risks down to a level comparable to human control. One of the reasons this can happen is if there are many possible anomaly situations that cannot be enumerated and tested for in advance (note, however, that if such situations occur in flight then human reactions are only slight less unpredictable than the responses of automated systems).

Human Control Unsafe: In some cases human control is unsafe. These include situations that require an extremely fast reaction time, cases in which humans are error-prone (e.g., because the task is boring or repetitive), or tasks that humans are simply not good at vs. automated systems.

DATA COLLECTION

Consider the automation of the task of monitoring and controlling one sub-system – the Space Shuttle Main Engine (SSME). Other automation tasks for other sub-systems have similar data collection issues. If we work through the decision diagram in figure 1 we generate the following questions:

- What reaction times are required? This data should be accessible from the engineering analysis of the engine. The primary complication is that some processes (e.g., controlling the firing of the engine) require quite fast response times, while other monitoring and control activities (e.g., long term trending for prognostics) require response times of days or weeks. Thus this step in the analysis may force the larger problem of “monitoring and control of SSME” to be broken into a series of sub-problems.
- What are the lifecycle costs of manual monitoring and control of SSME? Answering this question requires a detailed analysis of ground controller time plus the training time, support staffs, and other associated costs.
- What are the lifecycle costs of autonomous monitoring and control of SSME? This is one of the most difficult questions to answer because software development and validation costs are notoriously difficult to estimate. The best approach we know of is to develop two independent costing models: one based on a top-down decomposition of the work required, and a second based on analogy to past applications of autonomy (e.g., autonomous fault protection in unmanned planetary missions, prognosis for the Joint Strike Fighter, etc.).
- Do other barriers to using autonomy for monitoring and control of SSME exist? This is a more technical question. In the case of SSME, the outstanding question is whether technology for monitoring and control can handle the complex three-dimension chemical and physical interactions in SSME. Questions of this kind can be approximated by surveying the technical literature but are best answered definitely by building and testing prototype systems.

More generally, performing autonomy trades for CEV systems and sub-systems will require the following data sources:

- Detailed system and sub-system requirements. This is necessary to compute required reactions times and to estimate the expected load on the human crew if automated systems are not created.
- Lifecycle costs for past NASA missions. Ideally databases should be available showing how ground controller time is spent by task and by system and sub-system. Further, estimation methods need to

be devised to compute the training costs, facilities costs, and the costs to develop the procedures followed by ground and flight crews.

- Cost models for the development and validation of autonomous systems. As noted above, cost estimation for software is inexact in general and costs for autonomous systems are particularly hard to estimate. Research is needed on top down decomposition models to estimate development costs, the creation of databases of costs of past applications of autonomy, and accepted methods to cost and carry out validation of autonomous systems to the level required for human space flight. Further, detailed data should be kept going forward on the cost and schedule performance of projects building autonomous systems.
- Additional work is required on prototyping autonomous systems. The results of this work, both technically and programmatically (particularly the ability of the development teams to build the prototypes within schedule and budget estimates) should be captured in a standard format.

ISSUES IN AUTONOMY TRADE STUDIES

The complexities in performing autonomy trade studies center around three basic issues: the difficulty of ensuring that safety concerns are met, the difficulty of estimating the costs to develop and validate autonomy software and processes, and the difficulty of estimating the impacts of autonomy on operations costs. For the specific case of the CEV there is the additional challenge that the designs for the CEV, and the associated mission concepts of operation, continue to undergo rapid evolution.

Safety Concerns

Generally speaking, NASA is not willing to trade cost savings for reductions in safety margins. As mentioned above, autonomy software is a particularly complex case because some factors cause autonomous systems to be less safe than manual systems while other factors cause such systems to be safer. Faults in autonomous systems can cause unpredictable results that can in some cases be catastrophic. Further, human "common sense" knowledge may allow manual operation to be safer in the case of some kinds of unanticipated events. On the other hand, autonomous systems require that flight rules, and other operational constraints, are captured explicitly and are then not violated – avoiding failures that can be caused by manual procedures that do not always follow all flight rules. Further, humans are known to be error prone on some tasks, particularly repetitive tasks. Finally, autonomous systems can have much faster reaction times than manual systems – particularly if the manual system requires communication with earth.

On balance our approach to safety concerns is to allow safety concerns to dominate cost concerns. This means that if human control is less safe than automated control then our decision diagrams will favor automation. On the other hand, if automated control cannot be made as safe as human control then our decision diagrams will favor manual control. Further, in cases where either mode is fundamentally workable we assume that autonomy software would require sufficient validation to bring it up to at least the safety levels of human control.

Autonomy Development Costs

Development costs for software are often difficult to estimate. This is both because small changes in requirements can lead to large changes in development costs and because validation costs (and the rework required when validation fails) can vary widely. Autonomy software is particularly difficult to estimate precisely.

One interesting complication is that a major software cost driver is working around software errors discovered during operations. Autonomy software is generally harder to work around (precisely because humans are less involved in decision making). This reinforces the need to test and validate autonomy software during development. This need will increase development costs of autonomy software. On the other hand, upfront testing and validation will reduce the need for operational workarounds and thus decrease lifecycle costs (and increase safety during operations since workarounds are a common source of errors).

A similar complication arises during sustaining the engineering phase of a project. Certain approaches to autonomy (for example expert systems) effectively build the device model into the software. Thus non-trivial changes to the device require complete rework of the autonomous system. More recent approaches, particularly model-based planning and diagnosis, explicitly separate the device model from the reasoning engines. Our research group recently deployed such a model-based planner for activity planning on the MER rover. One positive result coming out of that effort was that we were able to support multiple flight rule and device model refinements with little or no change to the core autonomy code. Further, validation and verification was largely automated so redelivery of the software was a relatively straightforward process.

The best current approach to estimating software development costs is to combine a "bottom up" estimate that sums up all sub-tasks required for software development with a "top down" estimate that compares the proposed effort to past development projects. As we note below, deriving improved autonomy costing models is an important area for further work.

Autonomy Impacts on Operations Costs

As discussed above, the impact of autonomy on operations costs is not trivial to calculate. Consider a case where manual monitoring a sub-system requires a full time staff of 10 people. We cannot necessarily assume that an automated system would reduce headcount by 10. For example, some headcount will be needed to handle emergencies and to maintain the automated system. On the other hand, a monitoring staff of 10 may be supported by a larger "back office" staff for training, facilities, and process development. All such factors must be considered in estimating the impacts of autonomy. Further, experience with unmanned missions, and in the Defense Department, can be used to study analogous cases.

CEV Design Uncertainties

As yet there is no fixed design for the CEV. There are also no final mission scenarios. We do not know, for example, what the crew size will be, what the length of stay on the lunar surface will be, what power sources will be used, whether any sort of habitat construction will be attempted, what science tasks will be desired, etc. We clearly cannot wait for all design decisions to be finalized – both because this will take some time and because the level of autonomy will influence some of these decisions (e.g., crew size). Our current approach, described below, focuses on Shuttle and Station since these missions are somewhat similar to CEV spiral two. It should be understood, however, that this analysis will need to be adapted to account for the differences between Shuttle and Station and the final CEV designs.

CURRENT EFFORTS

We are currently engaged in an initial trade study on autonomy for CEV. The immediate purpose of this study is to perform a point analysis of the baseline practices in Shuttle and Station and then estimate the variance of bringing more automation into these processes. The results should be directly applicable to setting priorities for the second spiral of CEV development. More generally, this work is a primary step toward the creation of effective processes, tools, and models for performing ongoing autonomy trades in system and system-of-system design.

Specific tasks include:

- Perform a baseline analysis of the distribution of ground crew time for Station and Shuttle between the major systems. For each system document how time is spent (routine monitoring, failure isolation and recovery, crew support, etc.). (Year 1)
- Perform a variation analysis of how much time we can expect to save using commonly proposed process changes and by adding autonomous systems (e.g., mixed-initiative planning and scheduling, IVHM, automated failure detection, isolation, and recovery, interfaces for rapid situational awareness, and other techniques used in space science and by other government agencies such as DOD). Document how these savings are distributed between the ground crew and the flight crew. (Year 2)
- Estimate the expected costs to develop, and more importantly to flight validate, flight and ground autonomous systems. (Year 1)

- For each major system, analyze which autonomous systems can run on the ground and where drivers exist (for example, telemetry limitations or response time requirements) that will force the use of onboard autonomy. (Year 2)
- Roll these findings into recommendations for the distribution of control in each of the major CEV systems. (Year 2)

As discussed above, much of the data we would like to have for this trade study is unavailable. However, we will draw on several data sources:

- Flight controller time records. United Space Alliance maintains records of flight controller time to a low level of granularity. The challenge in this area will be in aggregating this data to a meaningful level. For example, to determine how much time automation of monitoring and control of a system would save it is necessary to aggregate the following: the console time spent monitoring the system, the time spent by "backroom" personnel supporting the controller, the time spent training controllers and other personnel, the time spent developing processes and procedures followed by controllers and other personnel, etc.

- Flight controller interviews. Johnson Space Center has conducted interviews with controllers as part of an attempt to asses the potential impact of autonomy on operations.

- Past experience with developing autonomy for flight missions. Relatively little autonomy has been used in manned missions. However, there are several autonomous systems that have been taken to TRL 9 for unmanned exploration. These include: the MER activity planner, the DS1 Remote Agent experiment, complex fault protection software developed for CASSINI, autonomous navigation and obstacle avoidance in MER, the automated planner developed by ARC for shuttle ground processing, and other systems. JPL and ARC have data on the costs and timelines of these development efforts.

- Past studies of operations concepts. For example, in 2000 JSC MOD (Mission Operations Division) created a missions operations concept for a mission to Mars that included the tasks to be allocated to ground, flight crew, and automation.

CONCLUSION AND SUGGESTED FUTURE WORK

Appropriate use of autonomy for the CEV will be a necessary part of NASA's efforts to minimize lifecycle costs. However, the required analysis is complex and, in many cases, the required data is either unavailable or non-existent. Efforts are ongoing to create the best possible set of recommendations based on available data. NASA's eventual goal should be to put autonomy trade studies on the same kind of formal footing that propulsion and other more mature technologies are on today. To this end we recommend that further work be undertaken to build reliable costing models both of autonomy software development and of the impact of autonomous systems on operations costs. Finally, we recommend that detailed databases of both development and operations costs be created and made widely available to the NASA community.

ACKNOWLEDGMENTS

The current trade study described here, "Trade Study on Autonomous Operations for the Crew Exploration Vehicle", is supported by the NASA Human and Robotic Technology initiative as part of the Advanced Studies, Concepts and Tools element. Many of the ideas described here were developed jointly with my colleges on that project: Charles Weisbin of the Jet Propulsion Laboratory, Cynthia Wells of United Space Alliance, Debra Schreckenghost of Metrica Inc., Mark Drummond of Ames Research Center, and Kanna Rajan of Ames Research Center.

REFERENCES

Chase, J.P., Elfes, A., Lincoln, W.P., Shishko, R.P., Rodriguez, G., Weisbin, C. "Identifying Technology Investments for Future Space Missions," in the proceedings of the *AIAA Space 2003 Conference*, Long Beach, CA, 2003.

Elfes, A., Lincoln, W.P., R.P., Rodriguez, G., Weisbin, Wertz, J.A. “A Risk-Based Framework for Optimization of Technology Portfolio Investment Choices for Early Space Mission Design,” *IEEE Transactions on Aerospace Engineering*, 2003.

Mackall, D., Nelson, S., and Schumann, J. “Verification and Validation of Neural Networks for Aerospace Applications,” Technical Report CR-211409, NASA, 2002.

Smith, J.H., Wertz, J., Weisbin, C. "Predicting the Cost of New Technology,” JPL Publication D26750, October 2003.

Smith, J.H., Wertz, J., Weisbin, C. "Building a Pathway to Mars: Technology Investment for Science return,” *Journal of Space Mission Architecture (JSMA),* **3**, pg 101, September 2003.

Weisbin, C., Lincoln, W., Elfes, A., and Rodriguez, G. "Relative Benefits of Potential Autonomy Technology Investments,” in the proceedings of the *International Conference on Space Mission Challenges for Information Technology (SMC-IT)*, Pasadena, California, 2003.